Exhibit 9

CHART FOR U.S. PATENT NO. 9,304,714 ("the '714 Patent")

Accused Products:

Dell's products, including but not limited to Dell's PowerFlex appliance and rack products (*e.g.*, PowerFlex R660, PowerFlex R760, PowerFlex R650, PowerFlex R750, PowerFlex R640, PowerFlex R740xd, PowerFlex R840) with PowerFlex software ("Accused Products"), infringe at least Claim 12 of the '714 Patent.

Claims	Exemplary Evidence of Infringement
12 [pre] A memory system, comprising:	To the extent the preamble is limiting, the Accused Products comprise a memory system. For example, the Accused Products comprise a "family of products . . . powered by PowerFlex software-defined storage," where "PowerFlex is storage virtualization software" and the "PowerFlex family . . . consists of a rack-level and two node-level offerings: an appliance and ready nodes." For example, the Accused Products "appl[y] the principles of server virtualization . . . creating . . . pools of block storage," "abstract[] the local storage contained within each server," and "pools all the storage resources together" into a "global pool." For example, the Accused Products comprise a "system [that] is the collection of entities managed by the Metadata Management (MDM) cluster." *See, e.g.*: PowerFlex software-defined infrastructure enables broad consolidation across the data center, encompassing almost any type of workload and deployment topology. Its software-first architecture enables automation and programmability of the complete infrastructure stack. It provides scalability, performance, and resiliency, enabling effortless adherence to stringent workload SLAs. As a universal infrastructure platform, PowerFlex combines compute and high-performance software-defined storage resources in a managed, unified fabric for both block and file. Available in flexible consumption options (rack, appliance, custom nodes, or in the public cloud), it enables various deployment architectures: independent compute and storage (two-layer), HCI (single-layer), or a mixture of the two. PowerFlex is ideal for high performance applications and databases, building an agile private/hybrid cloud, or consolidating resources in heterogeneous environments. To learn about the business value and benefits organizations have achieved by using PowerFlex to run and manage their important business workloads, please read this white paper.

Claims	Exemplary Evidence of Infringement
	System – A PowerFlex system is the collection of entities managed by the Metadata Management (MDM) cluster. **MDM** – Metadata Manager. A highly-available storage management cluster that resides alongside other software components within the system but sits outside the data path and supervises storage cluster health and configuration. It coordinates rebalancing and rebuilding/reprotecting data as changes occur in the system. **Protection Domain** – A protection domain is a logical entity that consists of a group of SDSs that provide data protection for each other. Each SDS belongs to one (and only one) protection domain. By definition, each protection domain is a unique set of SDSs. Protection domains can be added during installation and modified post-installation. **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. 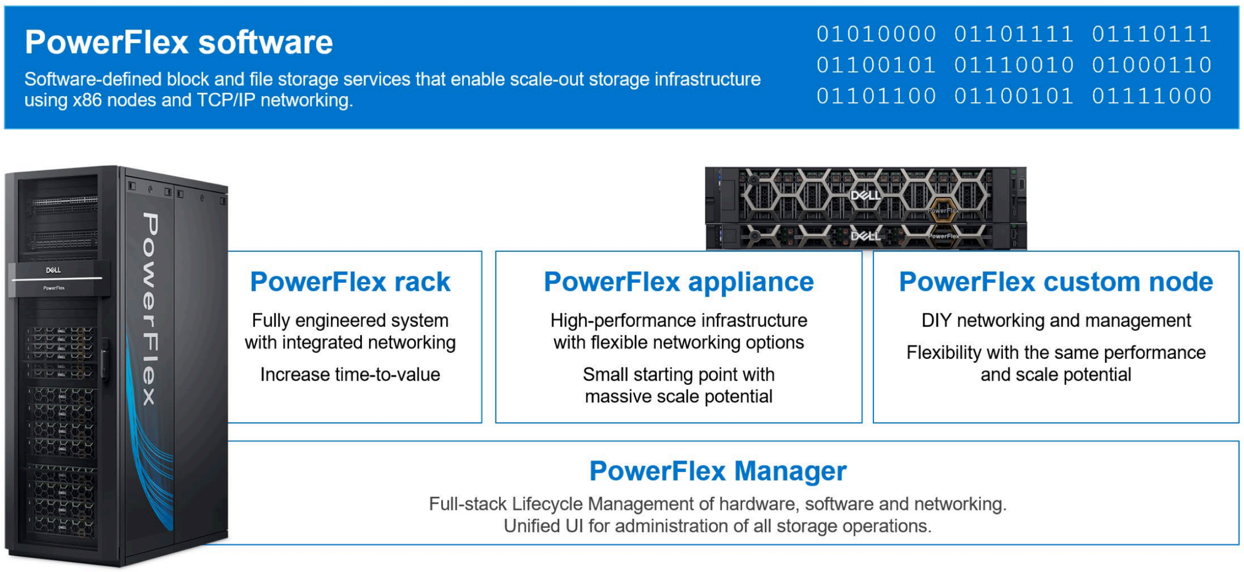

Claims	Exemplary Evidence of Infringement
	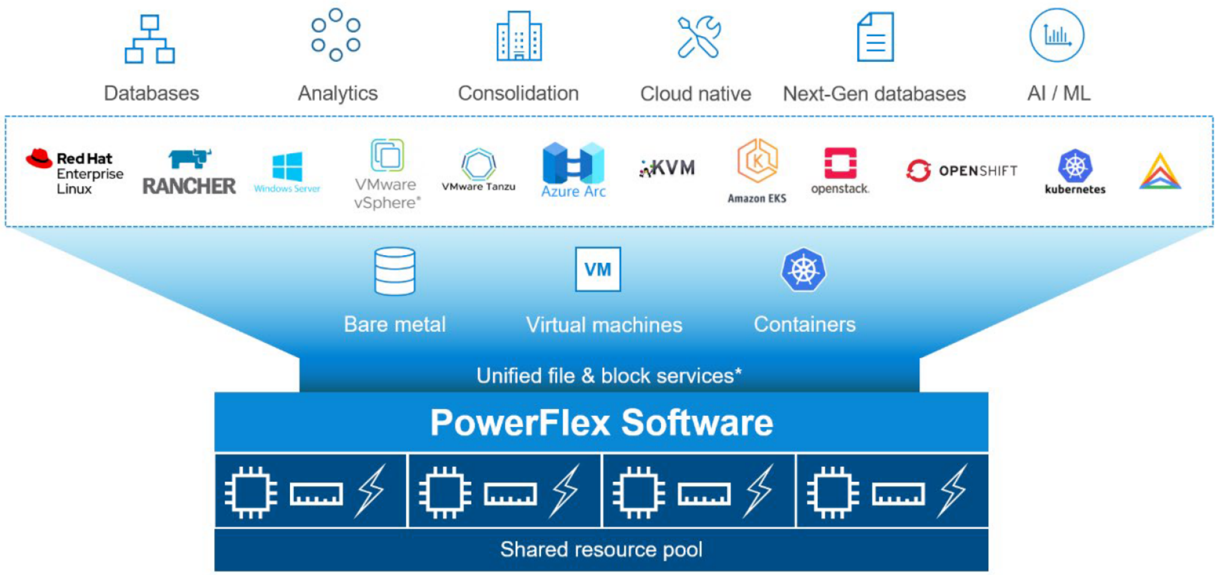 [Dell PowerFlex Specification Sheet] The Dell EMC™ PowerFlex™ family of products is powered by PowerFlex software-defined storage – a scale-out block storage service designed to deliver flexibility, elasticity, and simplicity with predictable high performance and resiliency at scale. Previously known as VxFlex OS, the PowerFlex storage software accommodates a wide variety of deployment options, with multiple OS and hypervisor capabilities.

Claims	Exemplary Evidence of Infringement
	PowerFlex software-defined storage architecture PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. 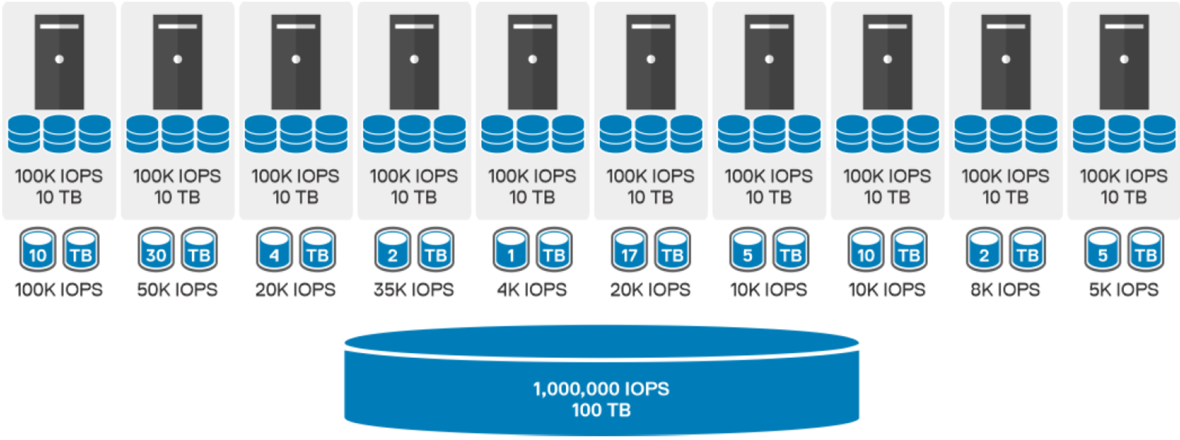 [Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview] The PowerFlex family currently consists of a rack-level and two node-level offerings: an appliance and ready nodes. This document primarily focuses on the storage virtualization software layer itself and is mostly relevant to the ready nodes, but it will be of interest to anyone wishing to understand the networking required for a successful PowerFlex-based storage system. PowerFlex is storage virtualization software that creates a server and IP-based SAN from direct-attached storage to deliver flexible and scalable performance and capacity on demand. As an alternative to a traditional SAN infrastructure, PowerFlex combines diverse storage media to create virtual pools of block storage with varying performance and data services options. PowerFlex provides enterprise-grade data protection, multi-tenant capabilities, and enterprise features such as inline compression, QoS, thin provisioning, snapshots and native asynchronous replication. PowerFlex provides the following benefits:

Claims	Exemplary Evidence of Infringement
	Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x
12 [a] a plurality of memory controllers in communication with a plurality of users;	The Accused Products comprise a plurality of memory controllers in communication with a plurality of users. For example, the Accused Products "fundamentally consists of three types of software components: the Storage Data Server (SDS), the Storage Data Client (SDC), and the Meta Data Manager (MDM)." For example, the "Storage Data Server (SDS) is a user space that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster" and comprises "a software service, running on a node that contributes disks to the storage cluster." For example, "[w]orking together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs" where a "storage pool is a set of physical storage devices," a "Device" is "[l]ocal, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool," and a "Volume" is "[a]nalogous to a LUN . . . a subset of a storage pool's capacity presented by an SDC as a local block device." For example, the "Storage Data Client (SDC) . . . [c]onsumes storage from the PowerFlex appliance," "provides front-end volume access to operating systems, applications, or hypervisors," presents "PowerFlex volumes as local block devices," "allows an operating system or hypervisor to access data served by PowerFlex clusters," "is a client-side software component," and is "analogous to a software HBA." For example, the "SDC allows shared volume access for uses such as clustering" but "does not require an iSCSI initiator, a fiber channel initiator, or an FCoE initiator." For example, the "Meta Data Manager[s]" or "MDMs control the behavior of the PowerFlex system . . . [t]hey determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue . . . directives to SDS components." *See, e.g.*: **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS** – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit.

Claims	Exemplary Evidence of Infringement
	SDC – Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands. **Device –** Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume –** Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet

Claims	Exemplary Evidence of Infringement
	PowerFlex fundamentally consists of three types of software components: the Storage Data Server (SDS), the Storage Data Client (SDC), and the Meta Data manager (MDM). Version 3.5 introduces a new component that enables replication, the Storage Data Replicator (SDR). 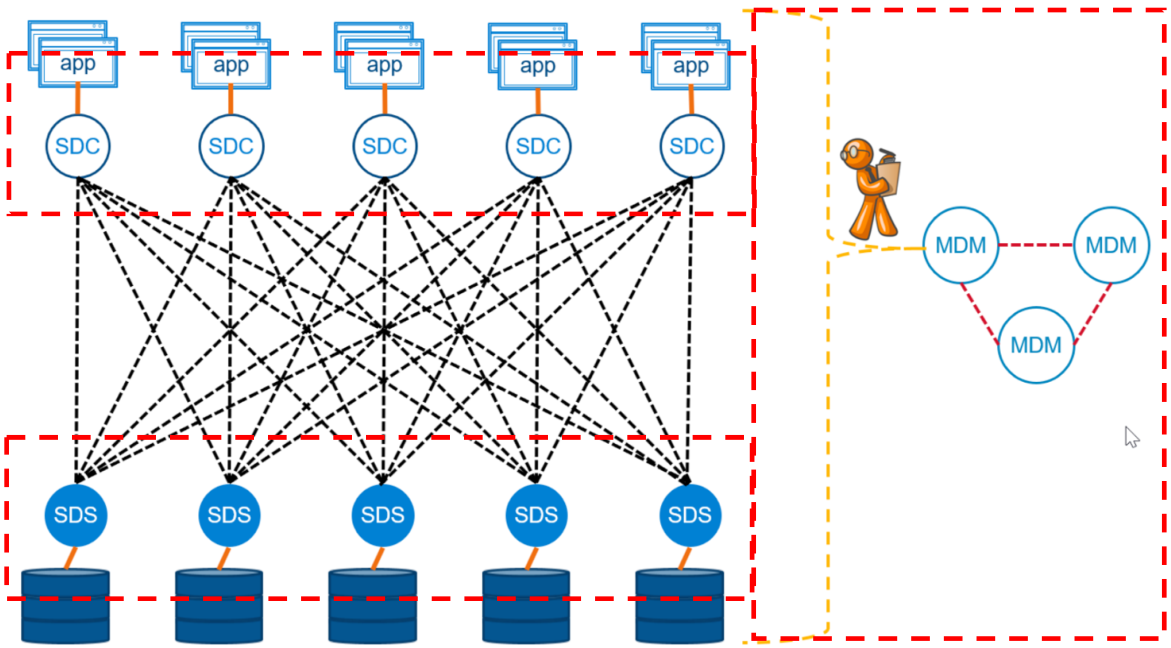 **Storage Data Server (SDS)** The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof.

Claims	Exemplary Evidence of Infringement
	SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. **Storage Data Client (SDC)** The Storage Data Client (SDC) allows an operating system or hypervisor to access data served by PowerFlex clusters. The SDC is a client-side software component that can run natively on Windows®, various flavors of Linux, IBM AIX®, ESXi® and others. It is analogous to a software HBA, but it is optimized to use multiple network paths and endpoints in parallel. The SDC provides the operating system or hypervisor running it with access to logical block devices called "volumes". A volume is analogous to a LUN in a traditional SAN. Each logical block device provides raw storage for a database or a file system and appears to the client node as a local device. The SDC knows which Storage Data Server (SDS) endpoints to contact based on block locations in a volume. The SDC consumes the distributed storage resources directly from other systems running PowerFlex. SDCs do not share a single protocol target or network end point with other SDCs. SDCs distribute load evenly and autonomously. The SDC allows shared volume access for uses such as clustering. The SDC does not require an iSCSI initiator, a fiber channel initiator, or an FCoE initiator. The SDC is optimized for simplicity, speed, and efficiency. A PowerFlex cluster may have up to 1024 SDCs. **Meta Data Manager (MDM)** MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Claims	Exemplary Evidence of Infringement
	PowerFlex runs on PowerFlex appliance nodes to operate the management and customer storage and tie in workloads. PowerFlex has the following components: • Storage data client (SDC): Consumes storage from the PowerFlex appliance • Storage data server (SDS): Contributes node storage to PowerFlex appliance • PowerFlex metadata manager (MDM): Manages the storage blocks and tracks data location across the system • Storage data replication (SDR): Enables replication on PowerFlex storage-only nodes Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. 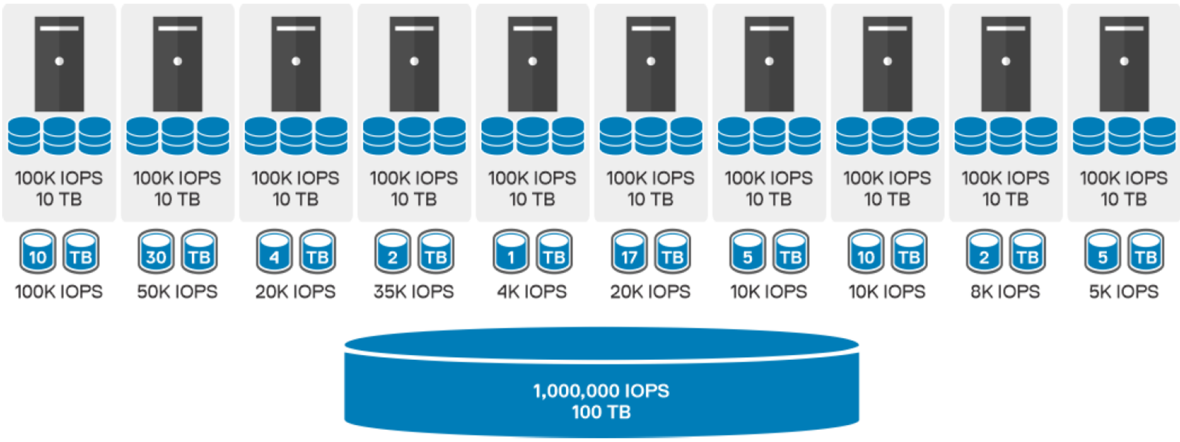 Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview
12 [b] a plurality of memory modules in communication with the	The Accused Products comprise a plurality of memory modules in communication with the plurality of memory controllers. For example, the Accused Products "fundamentally consist[ ] of three types of software components: the Storage Data Server (SDS), the Storage Data Client (SDC), and the Meta Data Manager (MDM)." For example, the "Storage Data Server (SDS) is a user space that aggregates raw local storage in a node and serves it out as part of a

Claims	Exemplary Evidence of Infringement
plurality of memory controllers;	PowerFlex cluster" and comprises "a software service, running on a node that contributes disks to the storage cluster." For example, "[w]orking together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs" where a "storage pool is a set of physical storage devices," a "Device" is "[l]ocal, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool," and a "Volume" is "[a]nalogous to a LUN . . . a subset of a storage pool's capacity presented by an SDC as a local block device." For example, the "Storage Data Client (SDC) . . . [c]onsumes storage from the PowerFlex appliance," "provides front-end volume access to operating systems, applications, or hypervisors," presents "PowerFlex volumes as local block devices," "allows an operating system or hypervisor to access data served by PowerFlex clusters," "is a client-side software component," and is "analogous to a software HBA." For example, the "SDC allows shared volume access for uses such as clustering" but "does not require an iSCSI initiator, a fiber channel initiator, or an FCoE initiator." For example, the "Meta Data Manager[s]" or "MDMs control the behavior of the PowerFlex system . . . [t]hey determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue . . . directives to SDS components." *See, e.g.*: **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS –** Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit. **SDC –** Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands. **Device –** Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume –** Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet

Claims	Exemplary Evidence of Infringement
	PowerFlex fundamentally consists of three types of software components: the Storage Data Server (SDS), the Storage Data Client (SDC), and the Meta Data manager (MDM). Version 3.5 introduces a new component that enables replication, the Storage Data Replicator (SDR). 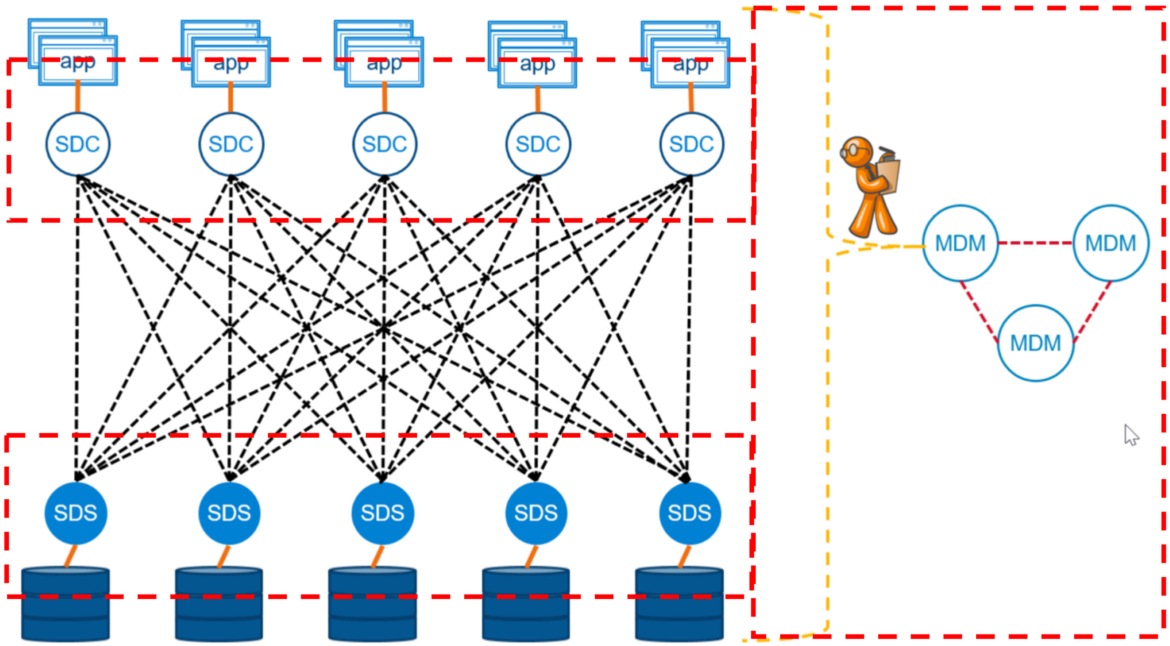 Storage Data Server (SDS) The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof.

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>
<h3>Storage Data Client (SDC)</h3>
The Storage Data Client (SDC) allows an operating system or hypervisor to access data served by PowerFlex clusters. The SDC is a client-side software component that can run natively on Windows®, various flavors of Linux, IBM AIX®, ESXi® and others. It is analogous to a software HBA, but it is optimized to use multiple network paths and endpoints in parallel.
The SDC allows shared volume access for uses such as clustering. The SDC does not require an iSCSI initiator, a fiber channel initiator, or an FCoE initiator. The SDC is optimized for simplicity, speed, and efficiency. A PowerFlex cluster may have up to 1024 SDCs.
<h3>Meta Data Manager (MDM)</h3>
MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components.
Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x
PowerFlex runs on PowerFlex appliance nodes to operate the management and customer storage and tie in workloads. PowerFlex has the following components:

Storage data client (SDC): Consumes storage from the PowerFlex appliance
Storage data server (SDS): Contributes node storage to PowerFlex appliance
PowerFlex metadata manager (MDM): Manages the storage blocks and tracks data location across the system
Storage data replication (SDR): Enables replication on PowerFlex storage-only nodes

Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide
</td></tr>
<tr><td>[c][i] wherein each controller of the plurality of controllers is operable to: maintain a pool of free memory areas of the plurality of</td><td>
Each controller in the Accused Products is operable to: maintain a pool of free memory areas of the plurality of memory modules at each controller of the plurality of controllers, wherein a logical unit is formed from free memory areas selected from the same pool of free memory areas.
For example, the Accused Products “appl[y] the principles of server virtualization to . . . servers with local disks, creating . . . shareable pools of block storage” and “abstract[] the local storage contained within each server,” “pool[ing] all the storage resources together” in a “global pool” of “resources . . . shared across the entire cluster.” For example, the Accused Products’ “SDS[s] work together [to] abstract local storage[and] maintain storage pools” where a “storage pool is a set of physical storage devices,” a “Device” is “[l]ocal, direct attached block
</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
memory modules at each controller of the plurality of controllers, wherein a logical unit is formed from free memory areas selected from the same pool of free memory areas;	storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool," and a "Volume" is "[a]nalogous to a LUN . . . a subset of a storage pool's capacity . . . ." For example, the Accused Products include "Meta Data Manager[s] (MDM[s])" that "determine and publish the mapping between clients and their volume data; they keep track of the system . . . ." For example, the "PowerFlex metadata manager (MDM) [m]anages the storage blocks and tracks data location across the system." For example, the Accused Products are aware of "allocated spare capacity and any generally free capacity" and can "reclaim unused space." *See, e.g.*: PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS** – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit. **SDC** – Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands. **Device** – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume** – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet

Claims	Exemplary Evidence of Infringement
	## Storage Data Server (SDS) The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. Protection from node, device, and network connectivity failure is managed with node-level granularity through **protection domains**. Protection domains are groups of SDSs in which user data replicas are maintained. **Fault sets** allow very large systems to tolerate multiple simultaneous node failures by preventing redundant copies from residing in a set of nodes (for example a whole rack) that might be likely to fail together. ## Meta Data Manager (MDM) MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x Protected maintenance mode makes the best use of all unused, available capacity, as it uses both the allocated spare capacity and any generally free capacity. It does not ignore capacity requirements. Nodes entering protected maintenance mode or in the same fault set may have degraded capacity. PowerFlex file systems provide increased flexibility by providing the ability to shrink and extend file systems as needed. Shrink and extend operations are used to resize the file system and update the capacity that is seen by the client. Extend operations do not change how much capacity is allocated to the file system. However, shrink operations may be able to reclaim unused space depending on how much capacity is allocated to the file system and the presence of snapshots or thin clones.

Claims	Exemplary Evidence of Infringement
	Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide Snapshots are a block image in the form of a storage volume or logical unit number (LUN) used to instantaneously capture the state of a volume at a specific point in time. Snapshots can be initiated manually or by new, automated snapshot policies. Snapshots in fine granularity storage pools are more space efficient and have better performance in comparison to medium granularity snapshots. PowerFlex supports snapshot policies based on a time retention mechanism. You can define up to 60 policy-managed snapshots per root volume A snapshot policy defines a cadence and the number of snapshots to keep at each level. Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview PowerFlex provides all this (and more) thanks to its "Secret Sauce" – its Distributed Mesh-Mirror Architecture. It ensures there are always two copies of your application data – thus ensuring availability in case of any hardware failure. Data is intelligently distributed across all the disk devices in each of the nodes within a storage pool. As more nodes are added, the overall performance increases nearly linearly, without affecting application latencies. Yet at the same time, adding more disks or nodes also makes rebuild times during those (admittedly rare) failure situations decrease – which means that PowerFlex heals itself more quickly as the system grows!

Claims	Exemplary Evidence of Infringement
	PowerFlex automatically ensures that the two copies of each block of data that gets written to the Storage Pool reside on different SDS (storage) nodes, because we need to be able to get a hold of the second copy of data if a disk or a storage node that holds the first block fails at any time. And because the data is written across all the disks in all the nodes within a Storage Pool, this allows for super-quick IO response times, because we access all data in parallel. Data also gets written to disk using very small chunk sizes – either 1MB or 4KB, depending on the Storage Pool type. Why is this? Doing this ensures that we always spread the data evenly across all the disk devices, automatically preventing performance hot-spots from ever being an issue in the first place. So, when a volume is assigned to a host or a VM, that data is already spread efficiently across all the disks in all Storage Nodes. For example, a 4-Node PowerFlex system, with 3 volumes provisioned from it, will look something like the following:  Figure 1: A Simplified View of a 4-Node PowerFlex System Presenting 3 Storage Volumes Now, here is where the magic begins. In the event of a drive failure, the PowerFlex rebuild process utilizes an efficient many-to-many scheme for very fast rebuilds. It uses ALL the devices in the storage pool for rebuild operations and will always rebalance the data in the pool automatically whenever new disks or nodes are added to the Storage Pool. This means that, as the system grows, performance increases linearly – which is great for future-proofing your infrastructure if you are not sure how your system will grow. But this also gives another benefit – as your system grows in size, *rebuilds get faster!*

Claims	Exemplary Evidence of Infringement
	Let me also highlight the "many-to-many" rebuild scheme used by each Storage Pool. This means that any data within a Storage Pool can be rebuilt to all the other disks in the same Pool. If we have 40 drives in our pool, it means that when one drive fails, the other 39 drives will be utilised to rebuild the data of the failed drive. This results in extremely quick rebuilds that occur in parallel, with minimum impact to application performance if we lose a disk: 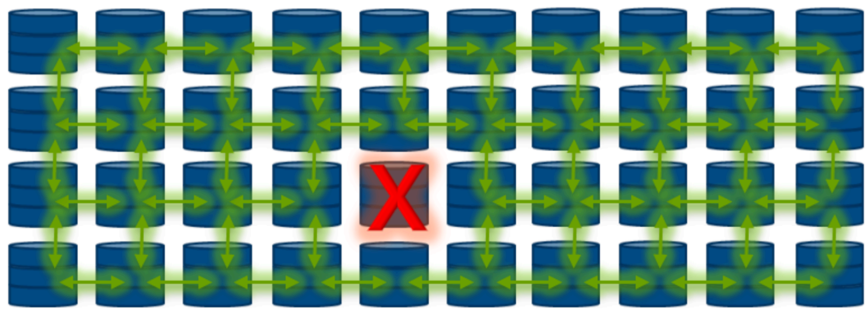 Figure 5: A 40-disk Storage Pool, with a Disk Failure… Showing The Magic of Parallel Rebuilds Resiliency Explained — Understanding the PowerFlex Self-Healing, Self-Balancing Architecture PowerFlex runs on PowerFlex appliance nodes to operate the management and customer storage and tie in workloads. PowerFlex has the following components: • Storage data client (SDC): Consumes storage from the PowerFlex appliance • Storage data server (SDS): Contributes node storage to PowerFlex appliance • PowerFlex metadata manager (MDM): Manages the storage blocks and tracks data location across the system • Storage data replication (SDR): Enables replication on PowerFlex storage-only nodes Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide
[c][ii] receive requests for maintenance of the logical unit from the user	The Accused Products receive requests for maintenance of the logical unit from the user at a first controller of the plurality of controllers, wherein the free memory areas are mapped to physical memory of the memory modules. For example, the Accused Products include "storage virtualization software" that "combine[] diverse storage media to create virtual pools of block storage" where a "storage pool is a set of physical storage devices" and "[a]nalogous

Claims	Exemplary Evidence of Infringement
at a first controller of the plurality of controllers, wherein the free memory areas are mapped to physical memory of the memory modules;	to a LUN, a volume is a subset of a storage pool's capacity presented to an SDC as a local block device." For example, in the Accused Products the "SDS . . . aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster" and "[a]cting together, SDSs maintain redundant copies of the user data." For example, "SDS components can communicate directly with each other, and collections of SDSs are fully meshed" where the "user data layout among SDS components is managed through storage pools . . . ." For example, "[t]raffic between the SDCs and SDSs" in the Accused Products "include all read and write traffic arriving at or originating from a client," where the "Storage Data Client (SDC)," a "client-side software component" that is "analogous to a software HBA" and "can run natively on Windows, various flavors of Linux, . . . and others allows an operating system or hypervisor to access data served by PowerFlex clusters." For example, the Accused Products' "SDC[s] provide[] . . . logical block addresses called 'volumes'" where "each logical block device provides raw storage for a database or a file system and appears to the client node as a local device," and "[c]lient volumes used by the SDCs are placed inside a storage pool." For example, the Accused Products' "Meta Data Manager[s] (MDM[s]) . . . control the behavior of the PowerFlex system," they "determine and publish the mapping between clients and their volume data [and] keep track of the state of the system." For example, the Accused Products "pool[] all the storage resources together" to create a "global pool," "[m]apping exposes the volume to the host, effectively creating a block device on the host," a "volume is distributed over all devices residing in the same storage pool," and those "resources are shared across the entire cluster." *See, e.g.*: PowerFlex is storage virtualization software that creates a server and IP-based SAN from direct-attached storage to deliver flexible and scalable performance and capacity on demand. As an alternative to a traditional SAN infrastructure, PowerFlex combines diverse storage media to create virtual pools of block storage with varying performance and data services options. PowerFlex provides enterprise-grade data protection, multi-tenant capabilities, and enterprise features such as inline compression, QoS, thin provisioning, snapshots and native asynchronous replication. PowerFlex provides the following benefits:

Claims	Exemplary Evidence of Infringement
	Storage Data Server (SDS) The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. Protection from node, device, and network connectivity failure is managed with node-level granularity through **protection domains**. Protection domains are groups of SDSs in which user data replicas are maintained. **Fault sets** allow very large systems to tolerate multiple simultaneous node failures by preventing redundant copies from residing in a set of nodes (for example a whole rack) that might be likely to fail together. **Meta Data Manager (MDM)** MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Claims	Exemplary Evidence of Infringement
	Storage Data Client (SDC) to Storage Data Server (SDS) Traffic between the SDCs and the SDSs forms the bulk of front-end storage traffic. Front-end storage traffic includes all read and write traffic arriving at or originating from a client. This network has a high throughput requirement. **Storage Data Server (SDS) to Storage Data Server (SDS)** Traffic between SDSs forms the bulk of back-end storage traffic. Back-end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, rebuild traffic, and volume migration traffic. This network has a high throughput requirement. **Traffic Types** PowerFlex performance, scalability, and security benefit when the network architecture reflects PowerFlex traffic patterns. This is particularly true in large PowerFlex deployments. The software components that make up PowerFlex (the SDCs, SDSs, MDMs and SDRs) converse with each other in predictable ways. **Architects designing a PowerFlex deployment should be aware of these traffic patterns in order to make informed choices about the network layout.** 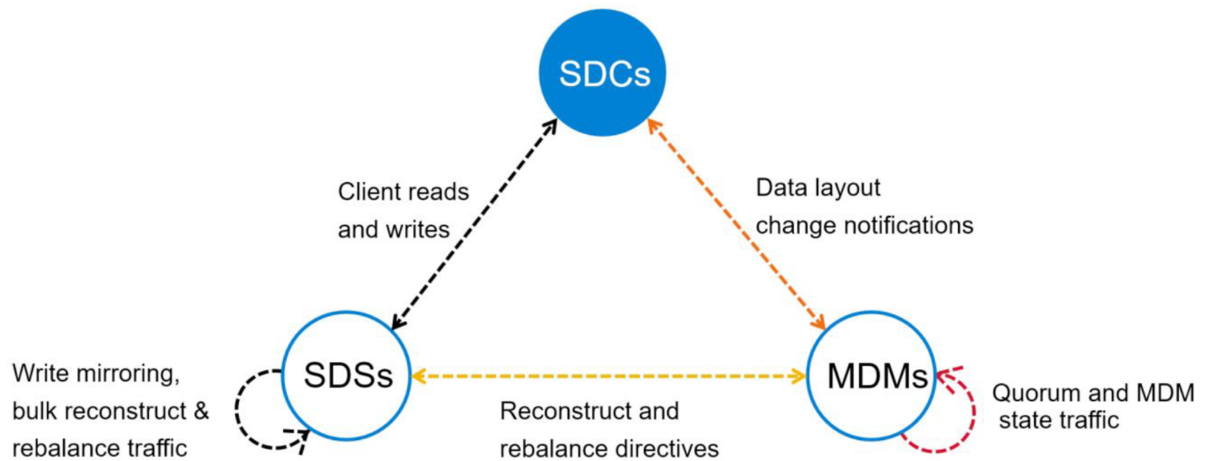

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Storage Pool - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool.

Device – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool.

Volume – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool.

Dell PowerFlex Specification Sheet

Mapping exposes the volume to the specified host, effectively creating a block device on the host. You can map a volume to one or more hosts.

A storage pool is a group of storage devices within a protection domain. Each time that you add devices to the system, you must map them to either storage pools or to acceleration pools. Create storage pools before you start adding SDSs and storage devices to the system.

A storage pool is a set of physical storage devices in a protection domain. A volume is distributed over all devices residing in the same storage pool. Add, modify, or remove a storage pool in the PowerFlex system.

Protected maintenance mode makes the best use of all unused, available capacity, as it uses both the allocated spare capacity and any generally free capacity. It does not ignore capacity requirements. Nodes entering protected maintenance mode or in the same fault set may have degraded capacity.

Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide

PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server.

PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster.

Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview</td></tr>
<tr><td>[c][iii] select a first free memory area from the pool</td><td>The Accused Products select a first free memory area from the pool of free memory areas and associate the first free memory area with the logical unit being maintained.</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
of free memory areas and associate the first free memory area with the logical unit being maintained;	For example, the Accused Products include "storage virtualization software" that "combine[] diverse storage media to create virtual pools of block storage" where a "storage pool is a set of physical storage devices" and "[a]nalogous to a LUN, a volume is a subset of a storage pool's capacity presented to an SDC as a local block device." For example, in the Accused Products the "SDS . . . aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster" and "[a]cting together, SDSs maintain redundant copies of the user data." For example, "SDS components can communicate directly with each other, and collections of SDSs are fully meshed" where the "user data layout among SDS components is managed through storage pools . . . ." For example, "[t]raffic between the SDCs and SDSs" in the Accused Products "include all read and write traffic arriving at or originating from a client," where the "Storage Data Client (SDC)," a "client-side software component" that is "analogous to software HBA" "can run natively on Windows, various flavors of Linux, . . . and others allows an operating system or hypervisor to access data served by PowerFlex clusters." For example, the Accused Products' "SDC[s] provide[] . . . logical block addresses called 'volumes'" where "each logical block device provides raw storage for a database or a file system and appears to the client node as a local device," and "[c]lient volumes used by the SDCs are placed inside a storage pool." For example, the Accused Products' "Meta Data Manager[s] (MDM[s]) . . . control the behavior of the PowerFlex system," they "determine and publish the mapping between clients and their volume data [and] keep track of the state of the system." For example, the Accused Products "pool[] all the storage resources together" to create a "global pool," "[m]apping exposes the volume to the host, effectively creating a block device on the host," a "volume is distributed over all devices residing in the same storage pool," and those "resources are shared across the entire cluster." *See, e.g.*: PowerFlex is storage virtualization software that creates a server and IP-based SAN from direct-attached storage to deliver flexible and scalable performance and capacity on demand. As an alternative to a traditional SAN infrastructure, PowerFlex combines diverse storage media to create virtual pools of block storage with varying performance and data services options. PowerFlex provides enterprise-grade data protection, multi-tenant capabilities, and enterprise features such as inline compression, QoS, thin provisioning, snapshots and native asynchronous replication. PowerFlex provides the following benefits:

Claims	Exemplary Evidence of Infringement
	## Storage Data Server (SDS) The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. Protection from node, device, and network connectivity failure is managed with node-level granularity through **protection domains**. Protection domains are groups of SDSs in which user data replicas are maintained. **Fault sets** allow very large systems to tolerate multiple simultaneous node failures by preventing redundant copies from residing in a set of nodes (for example a whole rack) that might be likely to fail together. ## Meta Data Manager (MDM) MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components.

Claims	Exemplary Evidence of Infringement
	Storage Data Client (SDC) to Storage Data Server (SDS) Traffic between the SDCs and the SDSs forms the bulk of front-end storage traffic. Front-end storage traffic includes all read and write traffic arriving at or originating from a client. This network has a high throughput requirement. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **Device** – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume** – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet Mapping exposes the volume to the specified host, effectively creating a block device on the host. You can map a volume to one or more hosts. A storage pool is a group of storage devices within a protection domain. Each time that you add devices to the system, you must map them to either storage pools or to acceleration pools. Create storage pools before you start adding SDSs and storage devices to the system. Protected maintenance mode makes the best use of all unused, available capacity, as it uses both the allocated spare capacity and any generally free capacity. It does not ignore capacity requirements. Nodes entering protected maintenance mode or in the same fault set may have degraded capacity. Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide

Claims	Exemplary Evidence of Infringement
	PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. 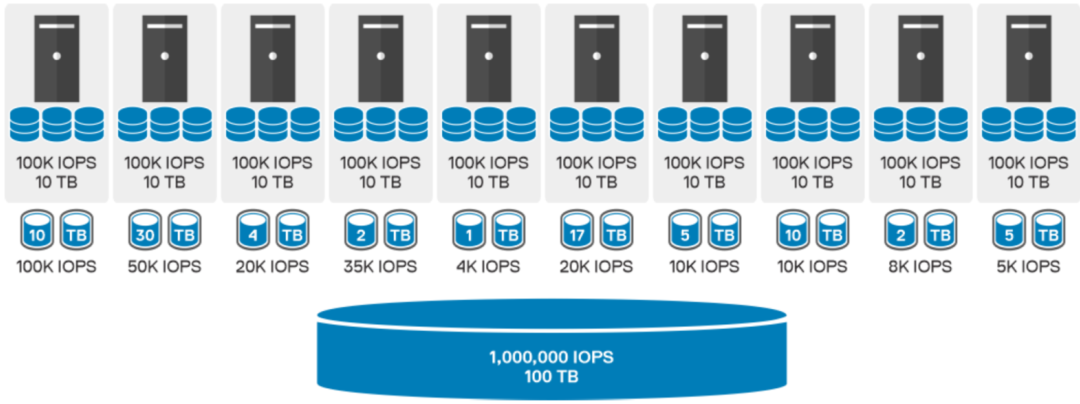 Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview
[c][iv] request each of the other controllers of the plurality of controllers to assign the first free memory area from the pool of free memory areas to the logical unit being maintained; and	The controller in the Accused Products request each of the other controllers of the plurality of controllers to assign the first free memory area from the pool of free memory areas to the logical unit being maintained. For example, the Accused Products "appl[y] the principles of server virtualization to . . . servers with local disks, creating . . . shareable pools of block storage" and "abstract[] the local storage contained within each server," "pool[ing] all the storage resources together" in a "global pool" of "resources . . . shared across the entire cluster." For example, "[w]orking together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs" where a "storage pool is a set of physical storage devices," a "Device" is "[l]ocal, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool," and a "Volume" is "[a]nalogous to a LUN . . . a subset of a storage pool's capacity presented by an SDC as a local block device." For example, the Accused Products include "Meta Data Manager[s] (MDM[s])" that "determine and publish the mapping between clients and their volume data; they keep track of the system . . . ." For example, the "PowerFlex metadata manager (MDM) [m]anages the storage blocks and tracks data location across the system." For example, the Accused Products "support[] two block protocols" where the "primary transport protocol is a proprietary TCP-based protocol that efficiently moves data between Storage Data Servers (SDSs) and Storage Data Clients (SDCs), as well as among the contributing SDSs." For example, the Accused Products "operate[] over an Ethernet fabric"

Claims	Exemplary Evidence of Infringement
	and while "many PowerFlex protocols are proprietary, all communications use TCP/IP transport." For example, "TCP allows . . . server and client pairs to exchange messages" and include "acknowledgement processes." For example, in the Accused Products "SDS components can communicate directly with each other, and collections of SDSs are fully meshed" and the "user data layout among SDS components is managed through storage pools, protection domains, and fault sets" where "[c]lient volumes used by the SDCs are placed inside a storage pool." *See, e.g.*: Data Access Protocols In addition to the file access protocols, listed above, PowerFlex supports two block protocols. The primary transport protocol is a **proprietary TCP-based protocol** that efficiently moves data between the Storage Data Servers (SDSs) and Storage Data Clients (SDCs), as well as among the contributing SDSs. The architecture includes native multipathing between the SDC and all SDSs that host volume data. The SDC translates this to a subset of the standard SCSI commands, for consumption by operating systems, hypervisors, and applications that can access raw block devices. **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS** – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit. **Device** – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume** – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet

Claims	Exemplary Evidence of Infringement
	PowerFlex operates over an Ethernet fabric. While many PowerFlex protocols are proprietary, all communications use standard TCP/IP transport. The following diagram provides a high-level overview of the port usage and communications among the PowerFlex software components. Some ports are fixed and may not be changed, while others are configurable and may be reassigned to a different port. For a full listing and categorization, see the "Port usage and change default ports" section of the Dell EMC PowerFlex Security Configuration Guide. 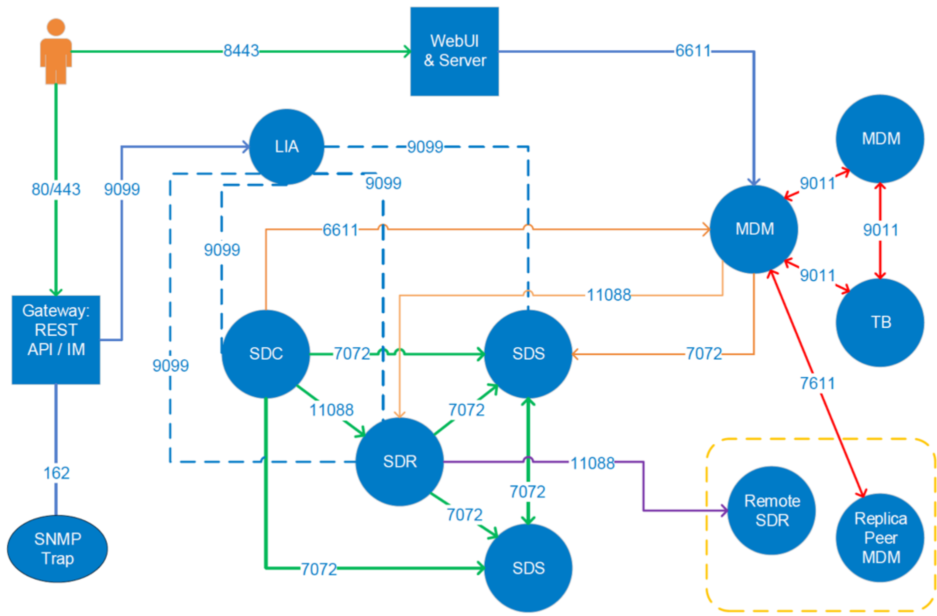

Claims	Exemplary Evidence of Infringement
	## Meta Data Manager (MDM) to Meta Data Manager (MDM) MDMs are used to coordinate operations inside the cluster. They issue directives to PowerFlex to rebalance, rebuild, and redirect traffic. They also coordinate Replication Consistency Groups, determine replication journal interval closures, and maintain metadata synchronization with PowerFlex replica-peer systems. MDMs are redundant and must continuously communicate with each other to establish quorum and maintain a shared understanding of data layout. MDMs do not carry or directly interfere with I/O traffic. The data exchanged among them is relatively lightweight, and MDMs do not require the same level of throughput required for SDS or SDC traffic. However, the MDMs have a very short (<400ms) timeout for their quorum exchanges, which happen every 100ms. **MDM to MDM traffic requires a stable, reliable, low latency network.** MDM to MDM traffic is considered back-end storage traffic. PowerFlex supports the use of one or more networks dedicated to traffic between MDMs. At a minimum, two 10 GbE links should be used per MDM for production environments, although 25GbE is more common. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x The Transmission Control Protocol (TCP) is a connection-oriented protocol used by the Internet Protocol (IP) transport layer. **Using a network connection, TCP allows a server and client pairs to exchange messages by using data segments packaged inside of data requests and responses.** In this tutorial, we'll learn how to initiate this connection by sending a Synchronize Sequence Numbers (SYN) packet. TCP provides a reliable connection for transferring data. Additionally, TCP has built-in error checking and guarantees the delivery of all data transmitted via its retransmission and acknowledgment processes. As a result, TCP is one of the most popular protocols in use on the internet. The protocol requires a connection. It is the responsibility of the client to make the initial request to start the connection. The control bits provided in the TCP protocol header facilitate this request to connect. SYN/ACK in the TCP Protocol

Claims	Exemplary Evidence of Infringement
	Storage Data Server (SDS) to Storage Data Server (SDS) Traffic between SDSs forms the bulk of back-end storage traffic. Back-end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, rebuild traffic, and volume migration traffic. This network has a high throughput requirement. **Storage Data Server (SDS)** The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. Protection from node, device, and network connectivity failure is managed with node-level granularity through **protection domains**. Protection domains are groups of SDSs in which user data replicas are maintained. **Fault sets** allow very large systems to tolerate multiple simultaneous node failures by preventing redundant copies from residing in a set of nodes (for example a whole rack) that might be likely to fail together. **Meta Data Manager (MDM)** MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Claims	Exemplary Evidence of Infringement
	PowerFlex runs on PowerFlex appliance nodes to operate the management and customer storage and tie in workloads. PowerFlex has the following components: • Storage data client (SDC): Consumes storage from the PowerFlex appliance • Storage data server (SDS): Contributes node storage to PowerFlex appliance • PowerFlex metadata manager (MDM): Manages the storage blocks and tracks data location across the system • Storage data replication (SDR): Enables replication on PowerFlex storage-only nodes PowerFlex enables flexible deployment options by allowing the separation of SDC and SDS components. PowerFlex Manager allows you to specify a non-root user instead of the root user when you configure a template for a compute-only, storage-only, hyperconverged, or PowerFlex file deployment. It addresses data center workload requirements through the following PowerFlex appliance deployment options: Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview
[c][v] receive a status or a success message from the other controllers	The Accused Products receive a status or a success message from the other controllers. For example, the Accused Products "appl[y] the principles of server virtualization to . . . servers with local disks, creating . . . shareable pools of block storage" and "abstract[] the local storage contained within each server," "pool[ing] all the storage resources together" in a "global pool" of "resources . . . shared across the entire cluster." For example, "[w]orking together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs" where a "storage pool is a set of physical storage devices," a "Device" is "[l]ocal, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool," and a "Volume" is "[a]nalogous to a LUN . . . a subset of a storage pool's capacity presented by an SDC as a local block device." For example, the Accused Products include "Meta Data Manager[s] (MDM[s])" that "determine and publish the mapping between clients and their volume data; they keep track of the system . . . ." For example, the "PowerFlex metadata manager (MDM) [m]anages the storage blocks and tracks data location across the system." For example, the Accused Products "support[] two block protocols" where the "primary transport protocol is a proprietary TCP-based protocol that efficiently moves data between Storage Data Servers (SDSs) and Storage Data Clients (SDCs), as well as among the contributing SDSs." For example, the Accused Products "operate[] over an Ethernet fabric" and while "many PowerFlex protocols are proprietary, all communications use TCP/IP transport." For example,

Claims	Exemplary Evidence of Infringement
	TCP allows . . . server and client pairs to exchange messages" and include "acknowledgement processes." For example, in the Accused Products "SDS components can communicate directly with each other, and collections of SDSs are fully meshed" and the "user data layout among SDS components is managed through storage pools, protection domains, and fault sets" where "[c]lient volumes used by the SDCs are placed inside a storage pool." *See, e.g.*: Data Access Protocols In addition to the file access protocols, listed above, PowerFlex supports two block protocols. The primary transport protocol is a **proprietary TCP-based protocol** that efficiently moves data between the Storage Data Servers (SDSs) and Storage Data Clients (SDCs), as well as among the contributing SDSs. The architecture includes native multipathing between the SDC and all SDSs that host volume data. The SDC translates this to a subset of the standard SCSI commands, for consumption by operating systems, hypervisors, and applications that can access raw block devices. **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS** – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit. **SDC** – Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands. **Device** – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume** – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet

4 PowerFlex TCP port usage

PowerFlex operates over an Ethernet fabric. While many PowerFlex protocols are proprietary, all communications use standard TCP/IP transport.

The following diagram provides a high-level overview of the port usage and communications among the PowerFlex software components. Some ports are fixed and may not be changed, while others are configurable and may be reassigned to a different port. For a full listing and categorization, see the "Port usage and change default ports" section of the Dell EMC PowerFlex Security Configuration Guide.

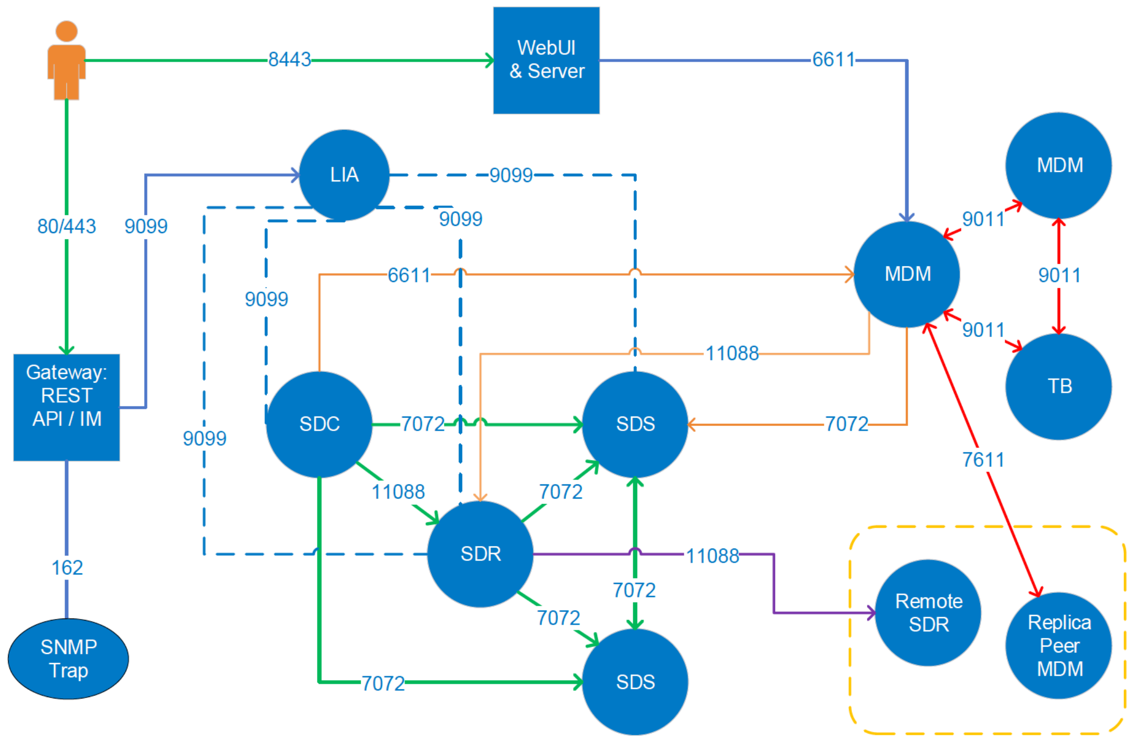

Figure 5 TCP port usage and communications within PowerFlex software-defined storage components. Arrows in the diagram indicate the direction of connection initiation. That is, the arrow points to a listening service port. Data may travel both directions over a connection after initiation. Dashed lines indicate that communication is internal to a node, among installed components.

Claims	Exemplary Evidence of Infringement
	Meta Data Manager (MDM) to Meta Data Manager (MDM) MDMs are used to coordinate operations inside the cluster. They issue directives to PowerFlex to rebalance, rebuild, and redirect traffic. They also coordinate Replication Consistency Groups, determine replication journal interval closures, and maintain metadata synchronization with PowerFlex replica-peer systems. MDMs are redundant and must continuously communicate with each other to establish quorum and maintain a shared understanding of data layout. MDMs do not carry or directly interfere with I/O traffic. The data exchanged among them is relatively lightweight, and MDMs do not require the same level of throughput required for SDS or SDC traffic. However, the MDMs have a very short (<400ms) timeout for their quorum exchanges, which happen every 100ms. **MDM to MDM traffic requires a stable, reliable, low latency network.** MDM to MDM traffic is considered back-end storage traffic. PowerFlex supports the use of one or more networks dedicated to traffic between MDMs. At a minimum, two 10 GbE links should be used per MDM for production environments, although 25GbE is more common. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x The Transmission Control Protocol (TCP) is a connection-oriented protocol used by the Internet Protocol (IP) transport layer. **Using a network connection, TCP allows a server and client pairs to exchange messages by using data segments packaged inside of data requests and responses.** In this tutorial, we'll learn how to initiate this connection by sending a Synchronize Sequence Numbers (SYN) packet. TCP provides a reliable connection for transferring data. Additionally, TCP has built-in error checking and guarantees the delivery of all data transmitted via its retransmission and acknowledgment processes. As a result, TCP is one of the most popular protocols in use on the internet. The protocol requires a connection. It is the responsibility of the client to make the initial request to start the connection. The control bits provided in the TCP protocol header facilitate this request to connect. SYN/ACK in the TCP Protocol

Claims	Exemplary Evidence of Infringement
	Storage Data Server (SDS) to Storage Data Server (SDS) Traffic between SDSs forms the bulk of back-end storage traffic. Back-end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, rebuild traffic, and volume migration traffic. This network has a high throughput requirement. **Storage Data Server (SDS)** The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance. Protection from node, device, and network connectivity failure is managed with node-level granularity through **protection domains**. Protection domains are groups of SDSs in which user data replicas are maintained. **Fault sets** allow very large systems to tolerate multiple simultaneous node failures by preventing redundant copies from residing in a set of nodes (for example a whole rack) that might be likely to fail together. **Meta Data Manager (MDM)** MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components.

Claims	Exemplary Evidence of Infringement
	Traffic Types PowerFlex performance, scalability, and security benefit when the network architecture reflects PowerFlex traffic patterns. This is particularly true in large PowerFlex deployments. The software components that make up PowerFlex (the SDCs, SDSs, MDMs and SDRs) converse with each other in predictable ways. **Architects designing a PowerFlex deployment should be aware of these traffic patterns in order to make informed choices about the network layout.** 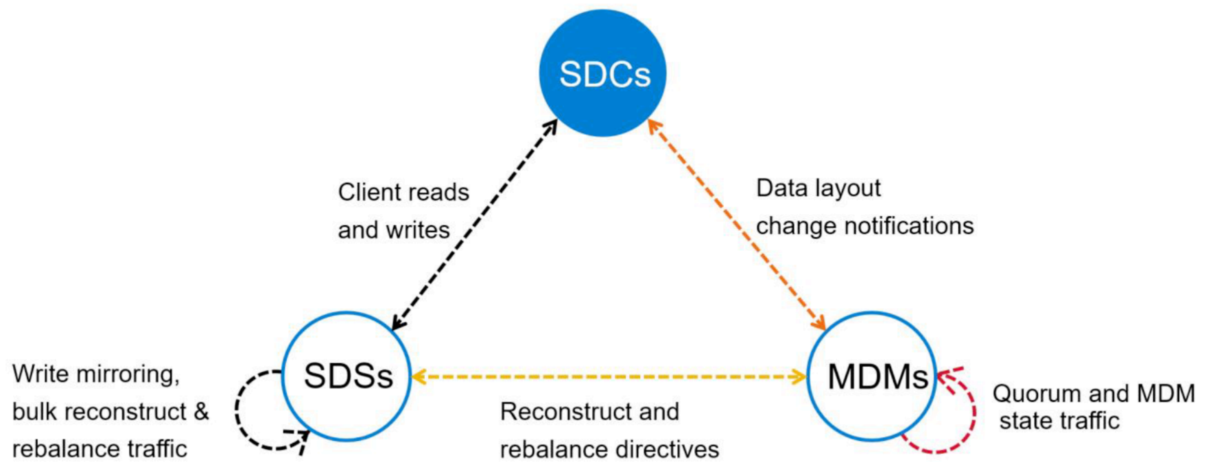 Figure 3 A simplified illustration of how the base PowerFlex software components communicate. A PowerFlex system will have many SDCs, SDSs, and MDMs. This illustration groups SDCs, SDSs, and MDMs. The arrows from the SDSs and MDMs pointing back to themselves represent communication to other SDSs and MDMs. Note that there is no SDC to SDC communication. The traffic patterns are the same regardless of the physical location of an SDC, SDS, or MDM. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>
PowerFlex runs on PowerFlex appliance nodes to operate the management and customer storage and tie in workloads. PowerFlex has the following components:

- Storage data client (SDC): Consumes storage from the PowerFlex appliance
- Storage data server (SDS): Contributes node storage to PowerFlex appliance
- PowerFlex metadata manager (MDM): Manages the storage blocks and tracks data location across the system
- Storage data replication (SDR): Enables replication on PowerFlex storage-only nodes

PowerFlex enables flexible deployment options by allowing the separation of SDC and SDS components. PowerFlex Manager allows you to specify a non-root user instead of the root user when you configure a template for a compute-only, storage-only, hyperconverged, or PowerFlex file deployment. It addresses data center workload requirements through the following PowerFlex appliance deployment options:

Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide

PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server.

PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster.

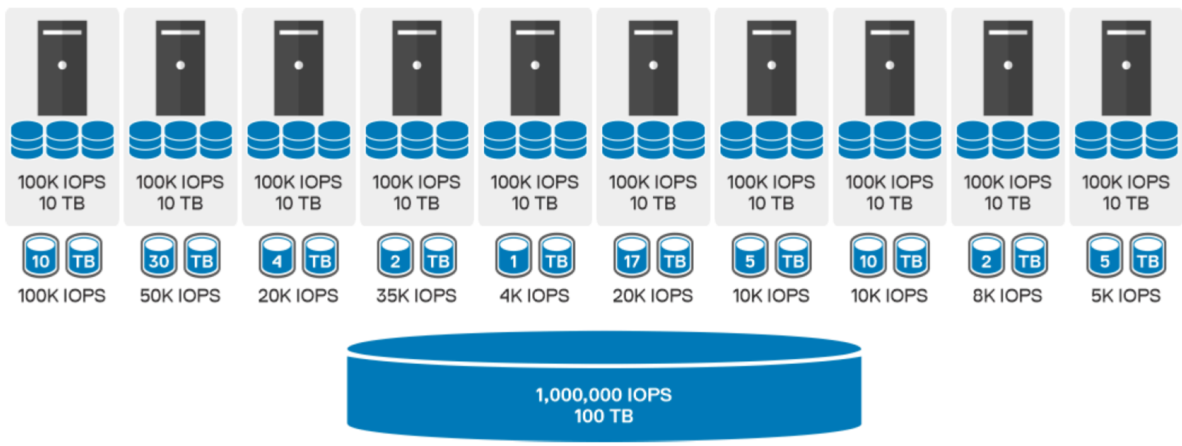

Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview
</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
[c][vi] wherein the first free memory area in the pool at each controller is associated with a same memory area of the memory modules.	The Accused Products associate the first free memory area in the pool at each controller with a same memory area of the memory modules. For example, the Accused Products include "storage virtualization software" that "combine[] diverse storage media to create virtual pools of block storage" where a "storage pool is a set of physical storage devices" and "[a]nalogous to a LUN, a volume is a subset of a storage pool's capacity presented to an SDC as a local block device." For example, in the Accused Products the "SDS . . . aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster" and "[a]cting together, SDSs maintain redundant copies of the user data." For example, "SDS components can communicate directly with each other, and collections of SDSs are fully meshed" where the "user data layout among SDS components is managed through storage pools . . . ." For example, "[t]raffic between the SDCs and SDSs" in the Accused Products "include all read and write traffic arriving at or originating from a client," where the "Storage Data Client (SDC)," a "client-side software component" that is "analogous to software HBA" and that can run natively on Windows, various flavors of Linux, . . . and others allows an operating system or hypervisor to access data served by PowerFlex clusters." For example, the Accused Products' "SDC[s] provide[] . . . logical block addresses called 'volumes'" where "each logical block device provides raw storage for a database or a file system and appears to the client node as a local device," and "[c]lient volumes used by the SDCs are placed inside a storage pool." For example, the Accused Products' "Meta Data Manager[s] (MDM[s]) . . . control the behavior of the PowerFlex system," "determine and publish the mapping between clients and their volume data [and] keep track of the state of the system." For example, the Accused Products "pool[] all the storage resources together" to create a "global pool," "[m]apping exposes the volume to the host, effectively creating a block device on the host," a "volume is distributed over all devices residing in the same storage pool," and those "resources are shared across the entire cluster." *See, e.g.*: PowerFlex is storage virtualization software that creates a server and IP-based SAN from direct-attached storage to deliver flexible and scalable performance and capacity on demand. As an alternative to a traditional SAN infrastructure, PowerFlex combines diverse storage media to create virtual pools of block storage with varying performance and data services options. PowerFlex provides enterprise-grade data protection, multi-tenant capabilities, and enterprise features such as inline compression, QoS, thin provisioning, snapshots and native asynchronous replication. PowerFlex provides the following benefits:

Claims	Exemplary Evidence of Infringement
	Storage Data Server (SDS) The Storage Data Server (SDS) is a user space service that aggregates raw local storage in a node and serves it out as part of a PowerFlex cluster. The SDS is the server-side software component. Any server that takes part in serving data to other nodes has an SDS service installed and running on it. A collection of SDSs form the PowerFlex persistence layer. Acting together, SDSs maintain redundant copies of the user data, protect each other from hardware loss, and reconstruct data protection when hardware components fail. SDSs may leverage SSDs, PCIe based flash, Storage Class Memory, spinning disk media, available RAM, or any combination thereof. SDS components can communicate directly with each other, and collections of SDSs are fully meshed. SDSs are optimized for rebuild, rebalance, and I/O parallelism. The user data layout among SDS components is managed through **storage pools, protection domains, and fault sets**. Client volumes used by the SDCs are placed inside a **storage pool**. Storage pools are used to logically aggregate similar types of storage media at drive-level granularity. Storage pools provide varying levels of storage service distinguished by capacity and performance.

Claims	Exemplary Evidence of Infringement
	Storage Data Client (SDC) The Storage Data Client (SDC) allows an operating system or hypervisor to access data served by PowerFlex clusters. The SDC is a client-side software component that can run natively on Windows®, various flavors of Linux, IBM AIX®, ESXi® and others. It is analogous to a software HBA, but it is optimized to use multiple network paths and endpoints in parallel. The SDC provides the operating system or hypervisor running it with access to logical block devices called “volumes”. A volume is analogous to a LUN in a traditional SAN. Each logical block device provides raw storage for a database or a file system and appears to the client node as a local device. The SDC knows which Storage Data Server (SDS) endpoints to contact based on block locations in a volume. The SDC consumes the distributed storage resources directly from other systems running PowerFlex. SDCs do not share a single protocol target or network end point with other SDCs. SDCs distribute load evenly and autonomously. The SDC is extremely lightweight. SDC to SDS communication is inherently multi-pathed across all SDS storage servers contributing to the storage pool. This stands in contrast to approaches like iSCSI, where multiple clients target a single protocol endpoint. The widely distributed character of SDC communications enables much better performance and scalability. The SDC allows shared volume access for uses such as clustering. The SDC does not require an iSCSI initiator, a fiber channel initiator, or an FCoE initiator. The SDC is optimized for simplicity, speed, and efficiency. A PowerFlex cluster may have up to 1024 SDCs. **Meta Data Manager (MDM)** MDMs control the behavior of the PowerFlex system. They determine and publish the mapping between clients and their volume data; they keep track of the state of the system; and they issue rebuild and rebalance directives to SDS components. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Claims	Exemplary Evidence of Infringement
	Storage Data Client (SDC) to Storage Data Server (SDS) Traffic between the SDCs and the SDSs forms the bulk of front-end storage traffic. Front-end storage traffic includes all read and write traffic arriving at or originating from a client. This network has a high throughput requirement. **Storage Data Server (SDS) to Storage Data Server (SDS)** Traffic between SDSs forms the bulk of back-end storage traffic. Back-end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, rebuild traffic, and volume migration traffic. This network has a high throughput requirement. **Storage Data Client (SDC) to Storage Data Replicator (SDR)** In cases where volumes are replicated, the normal SDC to SDS traffic is routed through the SDR. If a volume is placed into a Replication Consistency Group, the MDM adjusts the volume mapping presented to the SDC and directs the SDC to issue I/O operations to SDRs, which then pass it on to the relevant SDSs. The SDR appears to the SDC as if it were just another SDS. SDC to SDR traffic has a high throughput requirement and requires a reliable, low latency network. SDC to SDR traffic is considered front-end storage traffic. Dell EMC PowerFlex: Networking Best Practices and Design Considerations PowerFlex Version 3.5.x

Claims	Exemplary Evidence of Infringement
	System – A PowerFlex system is the collection of entities managed by the Metadata Management (MDM) cluster. **MDM** – Metadata Manager. A highly-available storage management cluster that resides alongside other software components within the system but sits outside the data path and supervises storage cluster health and configuration. It coordinates rebalancing and rebuilding/reprotecting data as changes occur in the system. **Protection Domain** – A protection domain is a logical entity that consists of a group of SDSs that provide data protection for each other. Each SDS belongs to one (and only one) protection domain. By definition, each protection domain is a unique set of SDSs. Protection domains can be added during installation and modified post-installation. **Storage Pool** - A storage pool is a set of physical storage devices within a protection domain. Each storage device belongs to one (and only one) storage pool. A volume is distributed over all devices residing in the same storage pool. **SDS** – Storage Data Server. A software service, running on a node that contributes disks to the storage cluster. Working together, several SDSs abstract local storage, maintain storage pools, and present volumes to the SDCs. Each SDS node is a fault unit, and the distributed mesh-mirror copies of data are never placed on the same fault unit. **SDC** – Storage Data Client. A client kernel driver that provides front-end volume access to operating systems, applications, or hypervisors. It presents PowerFlex volumes as local block devices. The SDC maintains peer-to-peer connections to every SDS managing a storage pool. It translates between the proprietary PowerFlex data transport protocol and block SCSI commands. **Device** – Local, direct attached block storage (DAS) in a node that is managed by an SDS and is contributed to a storage pool. **Volume** – Analogous to a LUN, a volume is a subset of a storage pool's capacity presented by an SDC as a local block device. A volume's data is evenly distributed across all disks comprising a storage pool, according to the data layout selected for that storage pool. Dell PowerFlex Specification Sheet Mapping exposes the volume to the specified host, effectively creating a block device on the host. You can map a volume to one or more hosts. A storage pool is a group of storage devices within a protection domain. Each time that you add devices to the system, you must map them to either storage pools or to acceleration pools. Create storage pools before you start adding SDSs and storage devices to the system. Storage Data Client (SDC) to Storage Data Server (SDS) Traffic between the SDCs and the SDSs forms the bulk of front-end storage traffic. Front-end storage traffic includes all read and write traffic arriving at or originating from a client. This network has a high throughput requirement. Storage Data Server (SDS) to Storage Data Server (SDS) Traffic between SDSs forms the bulk of back-end storage traffic. Back-end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, rebuild traffic, and volume migration traffic. This network has a high throughput requirement. Protected maintenance mode makes the best use of all unused, available capacity, as it uses both the allocated spare capacity and any generally free capacity. It does not ignore capacity requirements. Nodes entering protected maintenance mode or in the same fault set may have degraded capacity.

Claims	Exemplary Evidence of Infringement
	Dell PowerFlex Appliance with PowerFlex 4.x Administration Guide **PowerFlex software-defined storage architecture** PowerFlex applies the principles of server virtualization to standard x86 servers with local disks, creating high-performance, sharable pools of block storage. PowerFlex abstracts the local storage contained within each server. PowerFlex pools all the storage resources together. In the following figure, there is a global pool of 1 million IOPS and 100 terabytes, instead of having 100K IOPS and 10 terabytes available in each server. The applications are not constrained by what is within the local server, these resources are shared across the entire cluster. 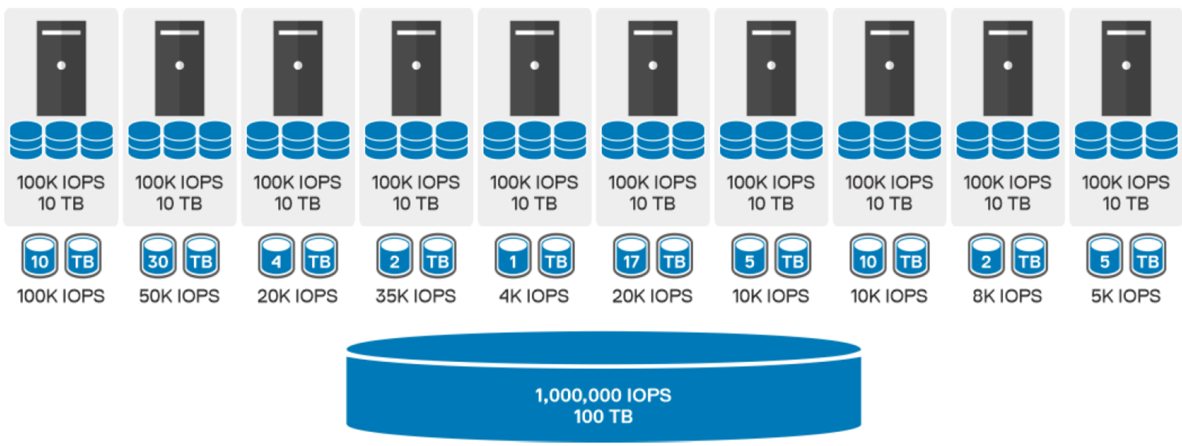 Dell PowerFlex Appliance with PowerFlex 4.x Architecture Overview